Controlling the Increasing Complexities of DB2 Data Management and Movement

DB2 managers eventually must deal with the critical issues of data availability and integrity in distributed production environments. Doing so successfully will require a firm foundation.

A number of fundamental changes are evolving in database management and data communications environments that will profoundly complicate the processes by which data is managed and moved. These changes are already being felt, especially in IBM's DB2 database management environment.

Information technology professionals preparing to deal with these challenges must make solution decisions that will protect the long-term investment of the business enterprise.

The changes appearing include:

- The inevitable growth of DB2
- The use of DB2 in production environments
- The complexities of distributed processing
- Movement toward an integrated MVS environment
- The emerging impact of downsizing

The depth of complexity that is beginning to appear will increase as organizations address the issues of movement and coordinated updating of data within a totally integrated MVS environment or among heterogeneous distributed databases.

The challenges become even more daunting when dealing with the issues of recovery and restart in these same environments. In short, these changes are creating a data management environment beyond the realm of people and procedures—one that demands the most advanced, reliable and synergistic software solutions available.

The inevitable growth of DB2

DB2 will continue to grow as the strategic database management system (DBMS) of choice. The demands of this growth will be manifested in terms of CPU consumption, DASD requirements, volume of data and number of users. Information Systems shops are moving rapidly from initial use of DB2, in which it is installed and employed primarily for development projects, to accelerating use, in which it is moved gradually into the production environment. As IS shops move finally to extensive use of DB2 (See chart, page 2), where it is used for high-volume, continuous-availability applications, they are quickly recognizing the hazards associated with managing performance, availability, integrity and control of data in this demanding arena.

The most obvious initial milestone is when DB2 is moved from a test or development stage to production. At this point, the need for high data availability and integrity begins to be felt by those who manage the DB2 subsystems. This proves true regardless of the size of the environment. Demands for high data availability and performance are suddenly felt in both 24-by-7 shops and those that have some form of batch window.

There are four major reasons for the inevitable growth of DB2. First, applications can be developed much faster with DB2 than with a non-relational DBMS. Once an organization realizes that the company saves money in development and gets an earlier use of the application, the trend toward use of DB2 for development accelerates. As it does, important data becomes a part of the new DB2-based applications. It would be inhibiting to the DB2 development process to insist that data remain positioned

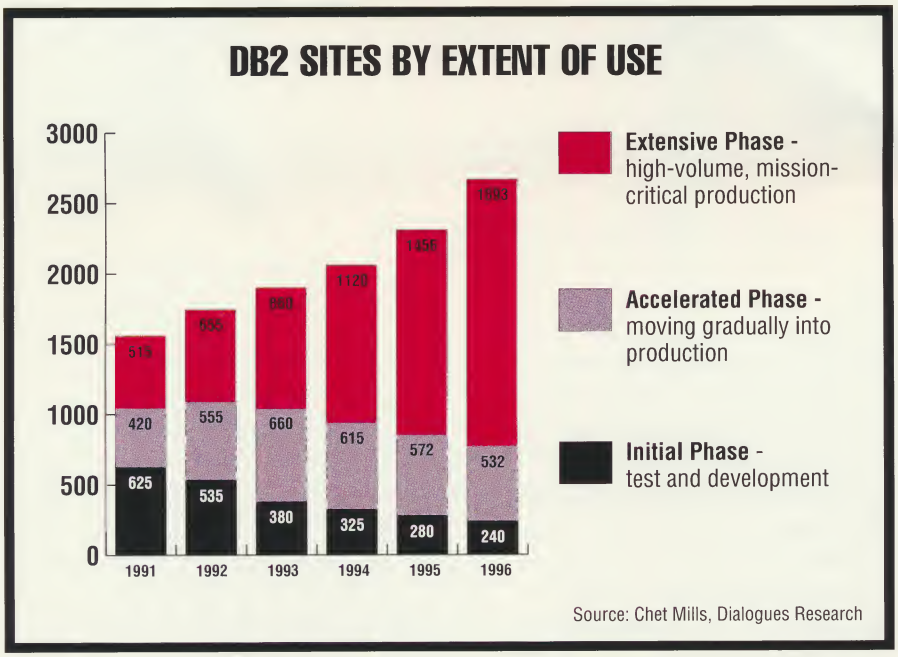

solely in a previous database management system. The result is a move toward increased repositioning of data in DB2 or toward a mixed application.

Second, applications are easier to maintain in DB2. Since most applications undergo significant maintenance over their life span, the ability to implement change more responsively improves the value of the application to the organization.

Third, the DB2 relational database management system is exceptionally capable in dealing with ad hoc requests for information without the need for program writing. This ability to respond quickly to informational needs often brings with it competitive advantage that is important to the business enterprise. As this becomes more obvious, the move to DB2 as the repository of choice for this type of data will accelerate.

Fourth, and most important, is that future technological advances—both software and hardware—will be built upon the DB2 foundation. Distributed database capabilities will evolve on a relational base. The relational implementation of DB2 is better able to take advantage of future hardware advances in memory and parallel processing. Major improvements are possible in the future through the synergy of relational concepts, operating system and hardware advances.

Other emerging technologies such as Information Warehouse, end-user computing capabilities, expert systems and object-oriented databases will be built from relational platforms.

Ironically, DB2's extensive capabilities are in themselves forcing changes in the way data must be managed. The many new uses available from DB2 are creating issues that DBAs have never had to face before. The result is that the role of the DBA is changing, moving toward that of a manager and facilitator and less of a controller. With that change has

The State of DB2 Today

by Richard Yevich, Senior Consultant, Codd & Date, Inc.

DB2, while still changing, has matured to the stage of large-scale, mission-critical applications. A database table of five billion rows, an online query system achieving close to 1,000 small transactions per second, and a batch system exceeding two million updates per hour in a 24-by-7 environment are in existence. In development, a major service firm is in the planning stages of a 130-billion-row table.

Data has matured into information, and DB2 is being selected as the information store. However, the issue of managing these data stores and achieving acceptable levels of performance requires more knowledge and use of tools than ever before.

The growth of DB2 use has far outpaced the growth of our understanding of how to manage, ensure performance and maintain integrity of the data. Complicating this further is the rapid movement into distributed systems, not only for EIS/DDS, but also for mission-critical OLTP systems.

The only downside to DB2 today is the lack of understanding of all the issues and automated tools required to manage this environment. DB2 will work exactly the way it is directed, GOOD OR BAD! With DB2, it is incredibly easy to implement poor systems and then not be able to understand, manage, monitor and correct the problems. It is a shame, because in the right environment, with the right tools, the flexibility of DB2 allows it to be corrected.

Richard Yevich is a consultant, software developer and lecturer with more than 23 years experience. He consults and lectures on relational, logical and physical database design, application design and performance issues.

come a need for fast and reliable data management tools.

Growth and new capabilities are placing new and complex demands on all data management resources—software, hardware and human.

DB2 and the critical production environment

As organizations increasingly extend DB2 usage into heavy-duty production environments, they are finding themselves faced with the same types of performance issues that exist with other DBMSs, but they are finding them much more difficult to control. DB2 data availability, integrity and recovery are now critical to the business itself, and the ability to manage data safely within the DB2 subsystem becomes vital.

To maintain the data and to keep DB2 performing efficiently, IS shops must schedule maintenance windows—periods during which they take the system down and perform maintenance jobs. These jobs include loading in new data, reorganizing table spaces, taking image copies and the like. Under these circumstances, time literally is money, and the speed and reliability of maintenance tasks becomes of paramount importance.

Most DB2 tuning and operations management must be done while the system is down, making data unavailable to users. Until recently, it was impossible to change performance control parameters (such as DSNZPARM values, buffer pool sizes, checkpoint frequency, timeout values or prefetch quantities or to change or repair an active log) without taking down the DB2 subsystem. To cancel runaway DB2 threads—a significant consumer of CPU —meant cancelling the originating address space and forcing the user off the system.

Software solutions will be necessary to ensure the high levels of performance and data availability vital to meet the demands of production environments.

Speed of recovery is another key area of concern for high-availability environments. Recovering a corrupted table space in DB2 by using the IBM recovery utility is labor-intensive and time-consuming. It involves collecting and merging data from image copies, incremental image copies, archive logs and active logs. If referential integrity is involved, not only the original table space, but also all those related to it must be recovered. For these reasons, it's impossible for an IS shop to accurately predict how long it will take to recover any given DB2 table space. The process can take a number of hours, to say the least, and the risk of error or data loss can be very high.

Getting Started With Information Warehouse

by Joyce Bischoff, President, Bischoff Consulting, Inc.

The good news about IBM's Information Warehouse is that it offers a framework for delivering business information across an enterprise. The bad news is that very few organizations are in a position to take advantage of it. The framework consists of enterprise data, data delivery, and decision support applications. Although organizations have a wealth of data, most of it is tightly bound to applications and lacks the flexibility and definition to support its movement into the Information Warehouse environment.

To take advantage of the Information Warehouse framework, a carefully planned data architecture should be built with the enterprise data model as a foundation. Without a well-planned architecture, much unnecessary DASD will be consumed by redundant data and the data structures will lack the flexibility to respond to changing user requirements. Since CPUs are not infinitely fast, considerable technical expertise will be required to create a physical implementation that performs well in a relational DBMS, such as DB2.

Strategies for data management and movement must also be planned. Data must be moved from the traditional production environment into the Information Warehouse. Tools and procedures must be in place to minimize resource down time and enable users to move data from the warehouse into personal databases on any CPU.

It is highly desirable to automate as much of the decision making as possible in the data maintenance area. Without properly maintained data, resource usage and response time will probably not be acceptable. If an organization builds an Information Warehouse, however, the rewards are potentially great: data will be a corporate asset that may be used to gain competitive advantage.

Joyce Bischoff has written articles and lectured internationally on DB2 design and performance issues and served as leader of the Task Force on DB2 Standards for Performance for GUIDE International.

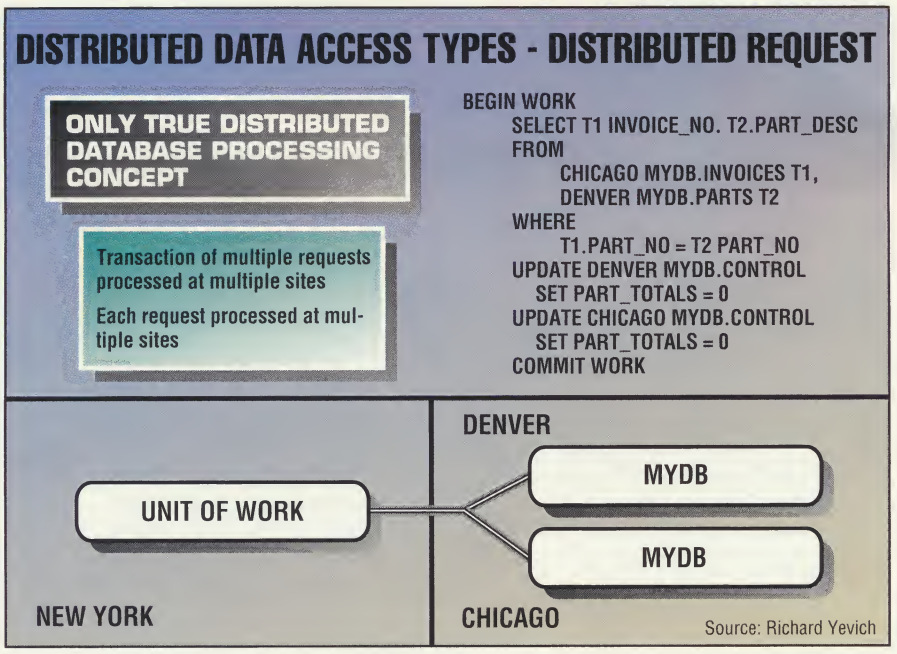

Distributed DB2: No simple solutions

Homogeneous distributed database implementation is already beginning to appear on a limited basis. True distributed access and distributed request capabilities in a heterogeneous environment will emerge as the predominant configuration of the future data management structure. As it does, DBAs are discovering, it will bring an exponential increase in data management complexities.

The very terms "distributed database" or "distributed processing" frequently mean different things to different people. For purposes of its Distributed Relational Database Architecture (DRDA), IBM defines remote request, remote unit of work, distributed request and distributed unit of work as four levels of distributed database activity.

Generally, when people speak of distributed database processing, they are referring to a distributed request, in which transactions of multiple requests are processed at multiple sites. This could be considered the ultimate goal of the truly distributed relational DBMS.

There are other, less advanced forms of distributed environments which, while not as complex as the distributed request scenario, still present significant challenges in terms of data availability, integrity, control and recovery.

It will be impossible to manage and move data in these complex environments without the use of highly sophisticated, well-designed software that will operate synergistically with other administrative tools and utilities.

The integrated MVS environment

Movement has already begun toward a totally integrated MVS environment in which data is managed and moved to and from DB2, other DBMSs and VSAM repositories throughout the VTAM network employing CICS, TSO or IMS/TM subsystems. This integrated data will be used not for development or testing but for high-volume, mission-critical applications programs that are vital to the business enterprise.

Although this may be achieved transparently to the various subsystems, the subsystems will be both logically and physically related and will have to be managed as a whole, because the applications force them to be.

If one subsystem suffers an outage, it affects the others. If one has to be recovered to a specific point in time, the others have to be recovered to that point in time. A single transaction has the power to trigger recoveries across all of the subsystems.

Solutions for this environment will have to be capable of managing data transparently across subsystem lines in a way that is non-disruptive to the normal functions of the subsystems.

The emerging impact of downsizing

Downsizing—the movement of corporate computing from the mainframe down to PCs and desktop work stations—is emerging as a significant factor in the way data will be managed and moved. Although originally perceived as a move away from the mainframe entirely, downsizing is now expected to evolve into a network that will typically include main-

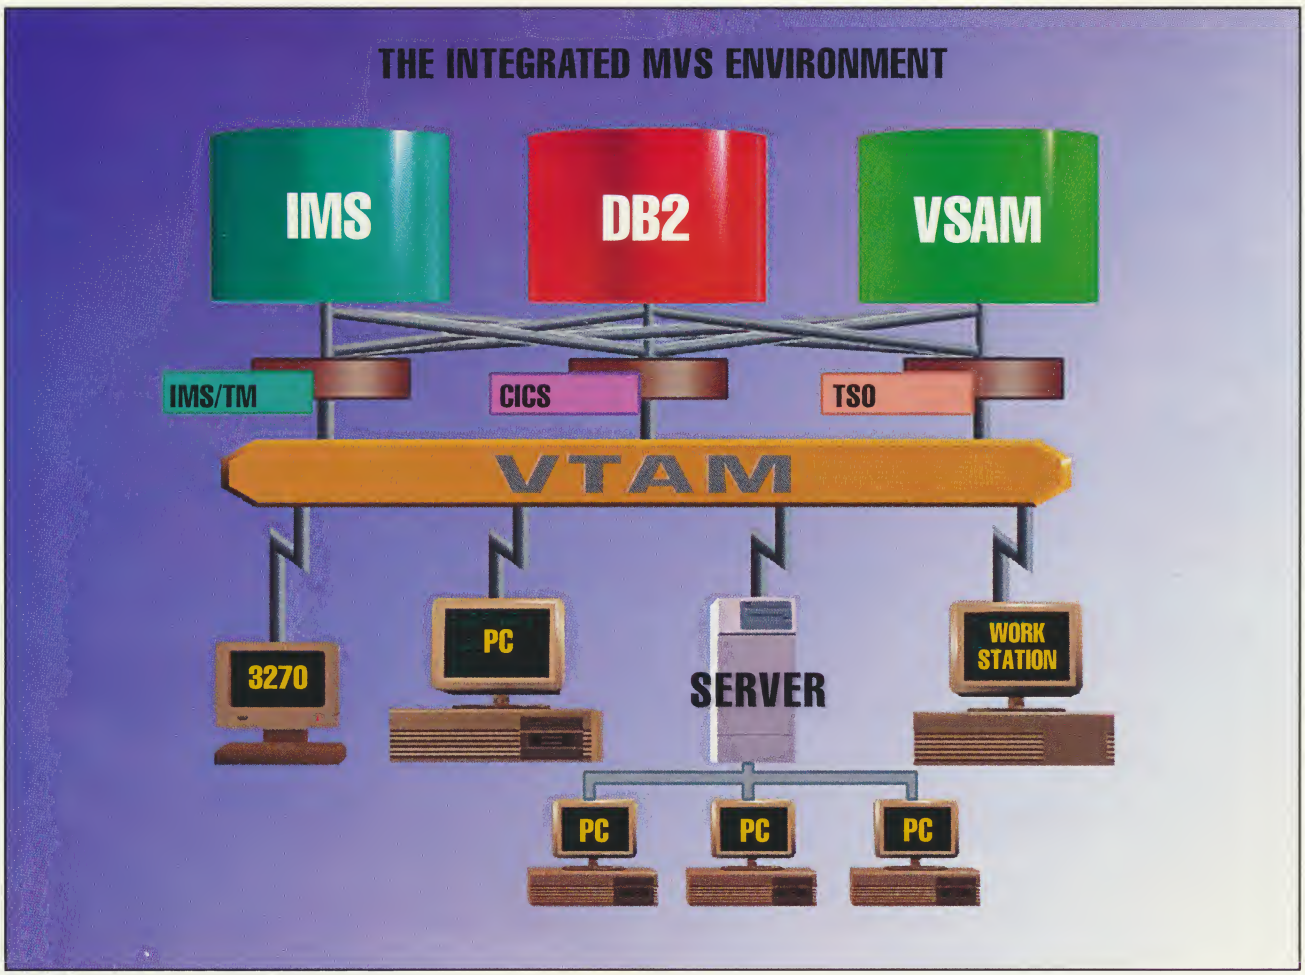

frames, servers and client work stations.

Downsizing offers two major advantages to corporations—lower hardware costs and faster end-user access to data. Businesses see downsizing as allowing the flow of business data out of corporate IS operations down to workgroup computers that can turn data into information to better manage the business. Rather than just seeking cheaper hardware, corporations are working toward making end users more productive.

Downsizing brings with it costs of its own. It requires significant investment in retraining—both for end users and programmers—and in new software licensing. The greatest challenge, however, is in managing the movement and control of data within the mainframe-client/server environment. Being able to do this in an automated way is extremely important, of course, because if the organization has to have an administrator on every LAN, that will effectively erase any savings achieved in hardware.

For the move to downsizing to be successful, organizations will need two sets of technically sophisticated tools — one for managing the environment and another for coordinating the extraction and propagation of corporate data within the network.

Old problems, new environment

The issues associated with cooperative processing and distributed databases will be basically the same as those that currently exist for any database management system, but they will be much more complex to resolve.

Database administrators and systems staff will be attempting to diagnose and correct performance problems in an environment in which any number of integrated DBMSs and data communications sub-

Downsizing and PC-Mainframe Coexistence

By Colin J. White
President, DataBase Associates

Downsizing has accelerated over the last few years primarily for two reasons. First, the price/performance of the PC compared to the mainframe has been quite attractive. Second, with the current shift in business and the economy, corporations need to be more competitive, and to be more competitive, they have to have speedy access to corporate data. Downsizing lets the PC access corporate data and turn it into information for use by business managers.

The word "downsizing" doesn't accurately describe what's occurring. I personally prefer the word "rightsizing." My perception is that there are applications that are ideally suited for work-group computing using networked PCs, but that the robustness of that environment today is not adequate to take over high-end corporate data processing applications. The so-called downsizing decision involves finding the right platform for each application. Rightsizing strategists see a need for networked PCs, but they will coexist in conjunction with mainframes. Consequently, corporations need to start looking at client-server and distributed computing to determine how desktop systems and work-group systems will work in conjunction with mainframes.

One of the key factors for successful downsizing will be managing these mainframe-to-PC environments. When organizations start developing and deploying applications that involve dozens, perhaps hundreds of PCs, the real challenge is going to be in managing and synchronizing the system software, applications software and PC configuration files. Unless we get well designed and automated tools in this area, the hardware cost savings of downsizing will be offset by the additional cost of administering such a system.

Another major hurdle is the movement of the data. About 80 to 85 percent of corporate data today is not in a relational form. It's in flat files, VSAM files and hierarchical databases, e.g., IMS. Secondly, a good deal of that data is not in a form that is usable to the end user. It's not well-structured for end-user access. In order to distribute the data to the PC level, organizations will have to copy that data to a relational information database and then stage it down to PCs.

For the move to downsizing to be successful, organizations will need tools to manage a distributed environment. Downsizing will require two sets of tools—one for managing the environment and another for coordinating the extraction and propagation of corporate data around the network.

Colin White is president of DataBase Associates, a database consulting firm based in Morgan Hill, California and is editor/publisher of InfoDB and Database Review. He has co-authored books on IBM's DB2 and SQL/DS products and is Conference Director for DB/EXPO.

systems could be at fault. Performance will be more difficult to monitor and much more difficult to tune.

When a transaction is distributed across multiple platforms, the chance of failure increases, and the risk is additive. For instance, to ensure an overall 99 percent availability of a transaction on five systems requires 99.8 percent availability on each platform. That is, what would have required a one-percent chance of failure on one platform must be reduced to a 0.2 percent chance of failure on each of the five platforms. In this case, the data manager has just lost 80 percent of his chance-of-failure pad. The need for availability is higher, the complexity of the environment is increasing, but the chance of failure must somehow be reduced.

Recovery to a point of consistency and a restart that preserves system-wide data integrity under these circumstances will present the most complex issues that information technology professionals have yet faced.

What the future will demand

Most organizations that have implemented DB2 rely on multiple DB2 subsystems to facilitate the design development, test and production stages of an application's life cycle. The data independence offered by DB2 allows the application data structures in each of these subsystems to vary for any given application.

These variations in structural elements are required to satisfy issues regarding performance, security and additional uses of data from one subsystem to another. Applications may also exist at various version levels on different subsystems. All of these factors make the task of database administration much more complex and error prone than with previous DBMS systems.

The strengths of DB2—data independence and flexibility—frustrate efforts to make implementing changes to DB2 structures a manageable process. The challenge faced by DB2 DBAs is to maximize the power of DB2's capabilities while keeping control of data structures and facilitating the responsive implementation of changes.

A key requirement for the information technology manager of the future will be a complete set of solutions for managing change within a complex DB2 environment that is comprised of multiple subsystems as well as CASE tools and various other repositories for structure definitions.

Change will have to be implemented while preserving structure modifications that have been made locally for tuning, data access or security purposes. CASE tools and other repositories used for application design must be synchronized with DB2's implementation of data structures. Bachman Information Systems, for example, provides such synchronization through its unique Model-Driven Development architecture.

Beyond the realm of people and procedures

The complexity continues to increase when one considers the level of effort required to recover and restart even a small part of any of these same operating environments. Recovery, of course, means with full data integrity and within a time frame that is acceptable to enterprise data availability needs. It becomes obvious that recovery from a system-wide failure that involves 10 DB2 subsystems residing in a heterogeneous distributed environment sharing data with 100,000 work stations where data also resides would be far too complex and time consuming for a people and procedures alone scenario.

It would have to be done through a software solution, and that software solution would require more development time, expense and sophistication than most IS shops and the vast majority of software vendors have or can afford.

It's a tough problem, but it's not impossible. The answers are to be found in technology leadership, quality, experience, and vision.

The need for technology leadership

To put it simply, it will be incumbent upon a third-party software developer to recognize these approaching problems and begin to work toward their solutions.

It will have to be an industry leader, one with extensive experience in the database management and data communications world. It should be one that has historically led the way with new and reliable technology. Finally, it must be one that is able to analyze and understand the complexity of the problems it has been asked to solve and one with the technological sophistication and capital strength to develop solutions to those problems.

In fact, that has already begun to happen.

One software vendor already provides an integrated set of high-performance tools that handles today's tactical problems and provides strategic solutions for the future.

One software vendor already has in place fast, reliable products that are horizontally integrated to provide solutions across multiple subsystems such as IMS, DB2, CICS, TSO and VTAM. These products fulfill day-to-day operating needs for data availability, integrity and recovery capabilities in complex integrated production environments.

One software vendor is already developing vertically integrated software solutions that will provide functional integration consistent with the strategic requirements of the future.

That vendor is BMC Software, Inc.

Why quality is critical

"There is no equivalent to a BMC product."

— Max Watson,
Chairman, President and
Chief Executive Officer,
BMC Software, Inc.

In the world of continuous-availability or 24-by-7 IS shops, there is no "almost" scenario. Database management systems and the administrative tools and utilities that support them must be fast and reliable. They must be designed to be safe and effective in the complex integrated environment of today and tomorrow. There are no bargain-basement answers.

The data management world continues to move into ever more complex integrated and distributed environments. The problems of maximizing performance, providing optimum data availability, protecting data integrity and ensuring database recovery and restart capabilities will demand a set of integrated, high-performance software solutions that span all of the subsystem environments.

For emerging DB2 shops, where initial concerns typically focus on development tools,

decisions must be made to protect the much larger investment in production. Selecting tools based on a single function will not create the solid foundation required for the future.

The only choice is BMC Software

BMC Software, Inc. has led the way in developing software solutions for IBM and IBM-compatible mainframe database management and data communications systems, including IMS, DB2, CICS and VTAM. BMC is the only software vendor that designs and develops solutions for data management and data movement across these integrated subsystems.

BMC is continuing to provide new products that meet the growing needs of this environment as it increases in size and complexity. It is for this purpose that BMC returns an industry-leading 23 percent of its revenue to research and development and product support—two areas in which BMC customers, in an independent survey, ranked BMC as the industry leader.

In addition, BMC is being called upon by concerned information technology professionals to provide solutions to the emerging problems of distributed processing—management of performance, change, recovery and restart. In fact, BMC began to address these issues some time ago as part of its research and development strategic planning. New products to meet these and other database management and data communications needs will be announced in the near future. Still others will follow.

The issues are complex. They require a robust foundation. Only BMC Software can offer the foundation and the demonstrated track record to deliver future products to meet the increasing complexities of data management and movement.

For reprints of this article or to talk to a BMC software consultant, call TOLL FREE in the U.S. and Canada, 1 800 841-2031 or FAX 713 242-6523.

BMC SOFTWARE

The Experience. The Technology. The Future.

P.O. Box 2002, Sugar Land, Texas 77487-2002 U.S.A.
Call toll free in the U.S. and Canada: 1 800 841-2031 • FAX: 713 242-6523